

PH	NL	MAT.
031272		DOSSIER
Europäisches Patentamt		
European Patent Office		
Office européen des brevets		



② Publication number:

0 395 271  
A2

②

### EUROPEAN PATENT APPLICATION

② Application number: 90303975.8

② Int. Cl. H04N 7/01, H04N 7/137,  
G06F 15/70, H04N 5/14

② Date of filing: 11.04.90

② Priority: 27.04.89 GB 8909651

② Applicant: SONY CORPORATION  
7-35 Kitashinagawa 6-Chome Shinagawa-ku  
Tokyo 141(JP)

② Date of publication of application:  
31.10.90 Bulletin 90/44

② Inventor: Samed, Raphael  
61 Farringdon Way, Tadley  
Basingstoke, Hampshire(GB)  
Inventor: Gillard, Clive Henry  
47 Saffron Close, Cufaude Village  
Chilham, Basingstoke, Hampshire(GB)


② Designated Contracting States:  
DE FR GB

② Representative: Thomas, Christopher Hugo et  
al  
D Young & Co 10 Staple Inn  
London WC1V 7RD(GB)

② Motion dependent video signal processing.

② In, for example, a motion compensated video standards converter wherein blocks in a first field or frame of a video signal are each compared with a plurality of blocks in the next succeeding field or frame of the video signal for deriving motion vectors representing the motion of the content of respective blocks between the first field or frame and the next field or frame, the blocks (a; q) adjacent to a border region of the field or frame are assigned motion vectors selected from the motion vectors which have been derived for those of the blocks (d, b, c; r, f, s, t, g) which are adjacent to the block (a, q) under consideration and are not in the border region.

FIG. 34.



Xerox Copy Centres

EP 0 395 271 A2

## MOTION DEPENDENT VIDEO SIGNAL PROCESSING

This invention relates to motion dependent video signal processing. More particularly, the invention relates to video standards converters using such signal processing, and to methods of deriving motion vectors representing motion between fields or frames of a video signal.

5 Video standards converters are well known devices used to convert video signals from one standard to another, for example, from a 625 lines per frame, 50 fields per second standard to a 525 lines per frame, 60 fields per second standard. Video standards conversion cannot be achieved satisfactorily merely by using simple linear interpolation techniques, because of the temporal and vertical alias which is present in a video signal. Thus, simple linear interpolation produces unwanted artefacts in the resulting picture, in particular, the pictures are blurred vertically and judder temporally.

10 To reduce these problems it has been proposed that video standards converters should use adaptive techniques to switch the parameters of a linear interpolator in dependence on the degree of movement in the picture represented by the incoming video signal.

It has also been proposed, for example for the purpose of data reduction in video signal processing, to generate motion vectors from an incoming video signal by a block matching technique, in which the content 15 of a search block in one field or frame is compared with the respective contents of a plurality of search blocks comprised in a search area in the following field or frame, to determine the minimum difference between the contents so compared, and hence the direction and distance of motion (if any) of the content of the original search block.

20 The present invention is particularly concerned with the problem of providing motion vectors for search blocks which are adjacent to a border region of the field or frame.

According to the present invention there is provided a motion compensated video standards converter comprising:  
means for comparing blocks in a first field or frame of a video signal with a plurality of blocks in the following field or frame of the video signal for deriving motion vectors representing the motion of the content 25 of respective said blocks between said first field or frame and said following field or frame;  
means to assign to each of said blocks adjacent to a border region of each said field or frame, motion vectors selected from motion vectors which have been derived for those of said blocks which are adjacent to said block under consideration and are not in said border region; and an interpolator controlled in dependence on said motion vectors.

30 According to the present invention there is also provided a method of deriving motion vectors representing motion between successive fields or frames of a video signal, the method including the steps of:  
comparing blocks in a first field or frame of the video signal with a plurality of blocks in the following field or frame of the video signal for deriving said motion vectors representing the motion of the content of 35 respective said blocks between said first field or frame and said following field or frame; and  
assigning to each of said blocks adjacent to a border region of each said field or frame, the motion vectors selected from the motion vectors which have been derived for those of said blocks which are adjacent to said block under consideration and are not in said border region.

The invention will now be described by way of example with reference to the accompanying drawings, 40 throughout which like parts are referred to by like references, and in which:

Figure 1 is a block diagram of an embodiment of motion compensated video standards converter according to the present invention;

Figure 2 shows diagrammatically progressive scan conversion;

45 Figures 3 to 6 show diagrammatically sequences of lines in sequences of fields for explaining progressive scan conversion;

Figure 7 is a block diagram showing the steps in motion adaptive progressive scan conversion;

Figure 8 shows diagrammatically progressive scanning, in particular the required estimate and difference value between successive fields;

Figures 9 and 10 are diagrams used in explaining the technique of Figure 8 in more detail, Figure 9 showing a progressive scan normalizing function and Figure 10 showing a progressive scan non-linear function;

Figure 11 shows diagrammatically the creation of pixels in missing lines in progressive scan conversion;

Figure 12 and 13 show diagrammatically search blocks and search areas, and the relationships therebetween;

Figure 14 shows a correlation surface;  
 Figures 15 and 16 show diagrammatically how a search block is grown;  
 Figure 17 shows the areas of a frame in which search block matching is not possible;  
 Figure 18 shows diagrammatically a moving object straddling three search blocks;  
 Figures 19 to 21 show three resulting correlation surfaces, respectively;  
 Figures 22 and 23 show further examples of correlation surfaces, used in describing a threshold test;  
 Figures 24 and 25 show still further examples of correlation surfaces, used in describing a rings test;  
 Figure 26 shows diagrammatically how the direction in which a search block is to grow is determined;  
 Figure 27 shows diagrammatically how a correlation surface is weighted;  
 Figure 28 shows the relationship between sample blocks and search blocks, and a frame of video;  
 Figure 29 shows motion vector regions in a frame of video;  
 Figures 30 to 32 show diagrams used in explaining motion vector reduction in respective regions of a frame of video;  
 Figure 33 shows part of the embodiment in more detailed block diagrammatic form;  
 Figure 34 is a diagram used in explaining motion vector assignment to search blocks which are adjacent to a border region of a frame;  
 Figure 35 is a flow chart of the operation of the arrangement of Figure 33;  
 Figures 36 and 37 show diagrammatically a first stage in motion vector selection;  
 Figures 38 and 39 show diagrammatically how a threshold is established during the motion vector selection;  
 Figure 40 shows diagrammatically a second stage in motion vector selection;  
 Figures 41 to 47 show arrays of pixels with associated motion vectors, used in explaining motion vector post-processing; and  
 Figure 48 shows diagrammatically the operation of an interpolator.

The embodiment of motion compensated video standards converter to be described is particularly intended for use in the conversion of a high definition video signal (HDVS) having 1125 lines per frame, 60 fields per second, to 24 frames per second 35mm film. However, it will be understood that the invention is not limited in this respect, and that the standards converter can readily be adapted to effect conversions between other standards.

Figure 1 is a block diagram of the standards converter. The standards converter comprises an input terminal 1 to which an input video signal is supplied. The input terminal is connected to a progressive scan converter 2 in which the input video fields are converted into video frames which are supplied to a direct block matcher 3 wherein correlation surfaces are created. These correlation surfaces are analysed by a motion vector estimator 4, which derives and supplies motion vectors to a motion vector reducer 5, wherein the number of motion vectors for each pixel is reduced, before they are supplied to a motion vector selector 6, which also receives an output from the progressive scan converter 2. Any irregularity in the selection of the motion vectors by the motion vector selector 6 is removed by a motion vector post processor 7, from which the processed motion vectors are supplied to and control an interpolator 8 which also receives an input from the progressive scan converter 2. The output of the interpolator 8, which is a standards-converted and motion-compensated video signal is supplied to an output terminal 9. Each part of the standards converter and the operation thereof will be described in more detail below.

The progressive scan converter 2 produces output frames at the same rate as the input fields. Thus, referring to Figure 2 which shows a sequence of consecutive lines in a sequence of consecutive fields, the crosses representing lines present in the input fields and the squares representing interpolated lines, each output frame will contain twice the number of lines as an input field, the lines alternating between lines from the input video signal and lines which have been interpolated by one of the methods to be described below. The interpolated lines can be regarded as an interpolated field of the opposite polarity to the input field, but in the same temporal position.

Progressive scan conversion is preferably carried out, for two main reasons; firstly, to make the following direct block matching process easier, and secondly in consideration of the final output video format. These two reasons will now be considered in more detail.

Direct block matching is used to obtain an accurate estimation of the horizontal and vertical motion between two successive video fields, as described in more detail below. However, due to the interleaved structure of the video signal on which direct block matching is performed, problems can arise.

Consider the image represented by Figure 3, which indicates a sequence of successive lines in a sequence of successive fields, the open squares representing white pixels, the black squares representing black pixels, and the hatched squares representing grey pixels. This, therefore, represents a static picture with a high vertical frequency component which in a HDVS would be 1125/3 cycles per picture height. As

this image has been sampled by the usual interface scanning procedure, each field appears to contain a static vertical frequency luminance component Y of 1125/6 cph, as indicated in Figure 4. However, the frequency components in each field are seen to be in anti-phase. Attempts to perform direct block matching between these two fields will lead to a number of different values for the vertical motion component, all of which are incorrect. This is indicated in Figure 5, in which the abbreviation LPF means lines per field. From Figure 5 it is clear that direct block matching will not give the correct answer for the vertical motion component, which component should in fact be zero. This is because the direct block matching is in fact tracking the alias component of the video signal rather than the actual motion.

Consider now Figure 6, which depicts the same static image as Figure 3, except that now each input field has been progressive scan converted to form a frame, the triangles representing Interpolated pixels. It can be seen that each frame now contains the same static vertical frequency component as the original input fields, that is 1125/3 cph. Thus, direct block matching between two successive frames can now give the correct value for the vertical motion, that is, zero, and the tracking of the vertical alias has been avoided. Moreover, there is the point that direct block matching on progressive scan converted frames will result in a more accurate vertical motion estimate, because the direct block matching is being performed on frames which have twice the number of lines.

Concerning consideration of the final output video format, in the case of the present embodiment, the converted video is supplied via tape to an electron beam recorder, and needs to consist of frames corresponding to the motion picture film rate of 24 frames per second. For this reason, therefore, the production of progressive scan converted frames is necessary, and moreover the progressive scan converted frames can also be used as a fall-back in the case where motion compensated standards conversion is deemed to be producing unacceptable results, for example, where the motion is too diverse to be analysed satisfactorily. In that case the use of the nearest progressive scan converted frame as the required output frame can produce reasonably acceptable results.

Progressive scan conversion can be carried out in a number of ways, such as by previous field replacement, median filtering in which three spatially consecutive lines are examined (temporally these three lines will come from two consecutive fields), or a motion compensated technique which utilizes multi-gradient motion detection followed by multi-direction linear interpolation. However, in the present embodiment the preferred method is motion adaptive progressive scan conversion, the steps of which are indicated in the block diagram of Figure 7. The concept is to use inter-field interpolation in wholly static picture areas to retain as much vertical information as possible, and to use intra-field interpolation when significant motion is present. This also aids smooth portrayal of motion. In scenes where the motion is somewhere between these two extremes, an estimate of the local motion present in the picture is made, and this is then used to mix together different proportions of inter- and intra-field interpolation.

In more detail, the modulus of the frame difference between previous and next fields is first generated, this being indicated in Figure 8. To generate the required estimates, the modulus inter-frame difference array from the previous and the next fields is generated at each point:

$$\Delta_U(\text{pixel, current line, current field}) = |Y(\text{pixel, current line, next field}) - Y(\text{pixel, current line, previous field})|$$

where:

$\Delta_U$  is the unnormalized modulus difference array, and  
 $Y$  is the luminance array corresponding to the 3D picture.

The modulus of difference is then normalized to adjust for the significance of changes in lower luminance areas:

$$\Delta_N(\text{pixel, current line, current field}) = F(\bar{Y}(\text{pixel, current line})) * \Delta_U(\text{pixel, current line, current field})$$

where:

$\Delta_N$  is the normalized modulus difference array

$\bar{Y}$  is the inter-frame average luminance value

$$Y(\text{pixel, current line}) = (Y(\text{pixel, current line, previous field}) + Y(\text{pixel, current line, next field})) / 2,$$

and  $F(\bar{Y})$  (the normalizing function) is derived as indicated in Figure 9.

The difference array  $\Delta$  is then vertically filtered together with the previous field difference by a three-tap filter (examples of coefficients are a quarter, a half, a quarter or zero, unity, zero) to reduce vertical alias problems, and in particular to minimize the problems encountered with temporal alias. Thus:

$$\Delta_F(\text{pixel, current line, current field}) = \Delta_N(\text{pixel, current line-1, previous field})C_1 + \Delta_N(\text{pixel, current line, current field})C_2 + \Delta_N(\text{pixel, current line+1, previous field})C_3$$

where:

$\Delta_F$  is the filtered normalized difference array, and

$C_1$  and  $C_2$  are filter coefficients, and  $2C_1 + C_2 = 1$  so that unity dc gain is maintained.

A vertical and horizontal intra-field filter of up to five taps by fifteen taps is then used to smooth the

difference values within the current field. In practice, a filter of three taps by three taps is satisfactory. Finally, in order to produce the actual motion estimation, a non-linear mapping function is applied using a function to provide the motion estimate (ME):

ME (pixel, current line) =  $\gamma$  (spatially filtered  $\Delta_t$  (pixel, current line))

5 The non-linear function  $\gamma$  is derived as shown in Figure 10, the static picture ME is zero, for full motion ME is one, and for intermediate motions a controlled transition occurs.

To produce an interpolated pixel, the pixels in the missing line are created by taking proportions of the surrounding lines as indicated in Figure 11. The motion estimate ME is then applied to the intra-frame interpolated value (generated from a two, four, six or preferably eight tap filter), and 1-ME is applied to the inter-field average (or alternatively to a more complex interpolated value), and these are summed to derive the progressive scan pixel estimate:

$$Y_{\text{out}}(\text{pixel, current line}) = \text{ME}(\text{pixel, current line}) * \left\{ \sum_{n=0}^{2} (Y_n(\text{pixel, current line-1-2n, current field}) + Y_n(\text{pixel, current line+1+2n, current field})) C_n \right\} + (1-\text{ME})(\text{pixel, current line}) * (Y_n(\text{pixel, current line, previous field}) + Y_n(\text{pixel, current line, next field})) Y_2$$

16 where:

$C_0, C_1, C_2$  and  $C_3$  are the intra-frame filter coefficients, and  $2(C_0 + C_1 + C_2 + C_3) = 1$  so that unity dc gain is maintained.

This method of progressive scan conversion is found to produce high quality frames from input fields, in particular because a moving object can be isolated and interpolated in a different manner to a stationary background.

17 Referring back to Figure 1, the frames of video derived by the progressive scan converter 2 are used to derive motion vectors. The estimation of motion vectors consists of two steps. Firstly, correlation surfaces are generated by correlating search blocks from consecutive frames. Then, having obtained these correlation surfaces, they have to be examined to determine the position or positions at which correlation is best. Several different methods of obtaining a correlation surface exist, the two main methods being phase correlation and direct block matching. There are, however, a number of problems associated with the use of phase correlation, these being very briefly problems relating to the transform mechanism, the windowing function, the block size and the variable quality of the contour of the surface produced. In the present embodiment, therefore, direct block matching is preferred.

30 The direct block matcher 3 operates as follows. Two blocks, respectively comprising a rectangular array of pixels from consecutive frames of the progressive scan converted video signal are correlated to produce a correlation surface from which a motion vector is derived.

35 Referring to Figure 12, firstly a small block called a search block of size 32 pixels by 23 lines is taken from a frame as shown in Figure 12. Then a larger block called a search area of size 128 pixels by 69 lines is taken from the next frame. The search block (SB) is then placed in each possible position in the search area (SA) as shown in Figure 13, and for each location the sum of the absolute difference of pixel luminance levels between the two blocks is calculated. This value is then used as the height of the correlation surface at the point at which it was derived. It can then be used in conjunction with other similarly derived values for each possible location of the search block in the search area to obtain a correlation surface, an example of which is shown in Figure 14. For clarity the surface is shown inverted, and as it is in fact the minimum that is required, the required point in Figure 14 is the main peak.

40 The size of the search block is selected by examining the minimum size of an object that may require motion compensation. For PAL 625 lines per frame, 50 fields per second signals a search block of 16 pixels by 8 lines has been found suitable for tracking a small object without allowing any surrounding information not within the object, but still within the search block, to affect the tracking of the object. This approach has therefore been adopted in the present embodiment, but modified to take account of the different numbers of active pixels per line, active lines per frame, and aspect ratio of a HDVS as compared with PAL 625/50. The comparative figures, the HDVS being put first, are as follows; 1920 (720) active pixels per line, 1035 (575) active lines per frame, 3:5.33 (3:4) aspect ratio.

45 It should be added that there is an argument for using a larger search block, since this means that a large object can be tracked. On the other hand, there exists an argument for using a smaller search block, to prevent a small object being overshadowed by the effect of a large object or background area. Also, however, there is the advantage that with small search blocks there is no requirement for the derivation of more than one motion vector from each of them. Because having a single motion vector is so much easier than having more than one, the present embodiment starts with a small search block as described above, and then causes the search block to grow into a bigger search block if no satisfactory result has been obtained. This then encompasses the advantages of both a small and a large search block. The criteria for a satisfactory result is set by the motion vector estimator 4 (Figure 1) referred to in more detail below and

which determines the motion vector from a given correlation surface.

This technique of causing the search block to grow is not only advantageous for tracking large objects. It can also help to track the movement of an object having the shape of a regular pattern of a periodic nature. Thus, consider Figure 15 where a search block A will match up with the search area B at locations 5 V1, V2 and V3, with each of them giving a seemingly correct measure of motion. In this case, however, the motion vector estimation, that is the process that actually analyses the correlation surface, will show that good correlation occurs in three locations which are collinear. The search block will therefore be caused to grow horizontally until it is three times its original width, this being the direction in which multiple correlation occurred in this case. The search area will also be correspondingly horizontally enlarged. As shown in Figure 16, with the enlarged search block 3A, there is only a single correlation point, which correctly relates to the motion of the object.

In this particular case the search block and the search area both have to grow horizontally, because the direction of multiple correlation is horizontal. It is equally possible, however, for the search block and the search area to grow vertically, or indeed in both directions, if the correlation surface suggests it.

15 It should be noted that block matching cannot be applied to all the search blocks in the frame, because in the border area there is not enough room from which a search area can be drawn. Thus, block matching cannot be effected in the border area of the frame shown hatched in Figure 17. This problem is dealt with by the motion vector reducer 5 (Figure 1) described in more detail below, which attempts to supply search blocks in this hatched area with appropriate motion vectors.

20 From the correlation surface (Figure 14) generated for each search block in a frame the motion vector estimator 4 (Figure 1) deduces the likely inter-frame motion between the search block and its corresponding search area. It should again be mentioned that for clarity all diagrams of correlation surfaces are shown inverted, that is, such that a minimum is shown as a peak.

The motion vector estimator 4 (Figure 1) uses motion vector estimation algorithms to detect the 25 minimum point on each correlation surface. This represents the point of maximum correlation between the search block and the search area, and hence indicates the probable motion between them. The displacement of this minimum on the correlation surface with respect to the origin, in this case the centre of the surface, is a direct measurement, in terms of pixels per frame, of the motion. For the simplest case, where the correlation surface contains a single, distinct minimum, the detection of the minimum point on the correlation surface is sufficient to determine accurately the motion between the search block and the search area. As previously mentioned, the use of small search blocks improves the detection of motion and the accuracy of motion estimation, but unfortunately small single search blocks are unable to detect motion in a number of circumstances which will now be described.

Figure 18 shows an object with motion vectors (5, 0) straddling three search blocks 1A, 2A and 3A in a 35 frame (t). When the search blocks 1A and 3A are correlated with respective search areas 1B and 3B in the next frame (t+1) a correlation surface shown in Figure 19 results showing a minimum at (5, 0). (This assumes a noiseless video source.) However, when the search block 2A is correlated with its respective search area 2B, the correlation surface shown in Figure 20 is produced, in which the search block 2A correlates with the search area 2B at every point in the y-axis direction. There is therefore no single 40 minimum in the correlation surface, and hence the motion between the search block 2A and the search area 2B cannot be determined.

However, now consider the situation if the search block 2A is grown such that it encompasses all three of the original search blocks 1A, 2A and 3A. When the grown search block 2A is correlated with a search area covering the original search areas 1B, 2B and 3B, the resulting correlation surface is as shown in Figure 21. This shows a single minimum at (5, 0) indicating the correct motion of the original search block 2A. This example illustrates the need for some unique feature in the source video, in order accurately to detect motion. Thus, the search blocks 1A and 3A both had unique vertical and horizontal features, that is the edges of the object, and hence motion could be determined. In contrast, the search block 2A had a unique vertical feature, but no unique horizontal feature, and hence horizontal motion could not be determined. However, by growing the search block until it encompasses a unique feature both horizontally and vertically, the complete motion for that search block can be determined. Moreover, it can be shown that growing the search block is beneficial when noise in the source video is considered.

55 A further example will now be considered with reference to Figure 22. This shows a correlation surface for a search block where the motion vector is (5, 3). However, due to the numerous other correlations which have taken place between the search block and the search area, the true motion is difficult to detect. An example of source video which might produce such a correlation surface would be a low contrast tree moving with the wind. It is now assumed that the search block and the search area are grown. The growing can take place in the horizontal direction, as in the previous example, or in the vertical direction, or in both

directions. Assuming that the neighbouring search blocks have the same motion, the mean effect on the resulting correlation surface will be to increase the magnitude of the minima at (5, 3) by a greater proportion than the magnitude of the other correlation peaks. This is shown in Figure 23, which indicates that it is then easier to detect the correct motion vector.

5 The way in which search blocks are grown will now be further considered with reference to Figure 18. Here it was required to grow the area of the search block 2A to encompass the areas of the search blocks 1A and 3A, and to produce the resulting correlation surface. In fact, the resulting correlation surfaces are produced directly by adding together the elements of the three correlation surfaces corresponding to the search blocks 1A, 2A and 3A. In effect, if each correlation surface is considered as a matrix of point 10 magnitudes, then the correlation surface of the enlarged search block 2A is the matrix addition of the correlation surfaces of the original search blocks 1A, 2A and 3A.

The area of the search block 2A could also be grown vertically by adding correlation surfaces of the search blocks above and below, whilst if the search block 2A is to be grown both horizontally and vertically, then the four neighbouring diagonal correlation surfaces have to be added as well. From this it will be seen 15 that the actual process of growing a search block to encompass neighbouring search blocks is relatively easy, the more difficult process being to decide when growing should take place, and which neighbouring search blocks should be encompassed. Basically, the answer is that the area of the search blocks should be grown until a good minimum or good motion vector is detected. It is therefore necessary to specify when 20 a motion vector can be taken to be a good motion vector, and this can in fact be deduced from the examples given above.

In the example described with reference to Figures 18 to 21, it was necessary to grow the search block horizontally in order to encompass a unique horizontal feature of the object, and hence obtain a single minimum. This situation was characterized by a row of identical minima on the correlation surface of Figure 20, and a single minimum on the correlation surface of Figure 21. From this the first criteria for a good 25 minimum can be obtained; a good minimum is the point of smallest magnitude on the correlation surface for which the difference between it and the magnitude of the next smallest point exceeds a given value. This given value is known as the threshold value, and hence this test is referred to herein as the threshold test.

It should be noted that the next smallest point is prevented from originating from within the bounds of a 30 further test, described below, and referred to herein as the rings test. In the case of a rings test employing three rings, the next smallest point is prevented from originating from a point within three pixels of the point in question. In the example of Figures 18 to 21, the correlation surface of Figure 20 would have failed the threshold test; the search area 2A is therefore grown and, given a suitable threshold value, the correlation surface of Figure 21 will pass the threshold test.

35 The threshold test can also be used to cause growing in the example described above with reference to Figures 22 and 23. Prior to growing the search block, the correct minimum is undetectable, due to the closely similar magnitudes of the surrounding points. Given a suitable threshold value, however, the correlation surface will fail the threshold test, and the search block will be grown. As a result, it will then be possible to detect the minimum among the other spurious points.

40 It will be seen that the use of a threshold is a subjective test, but the correct threshold for the correlation surface under test can be selected by normalizing the threshold as a fraction of the range of magnitudes within the correlation surface. This also lessens the effect of, for example, the contrast of the video source.

The rings test, referred to briefly above, and which is far less subjective, will now be further described. 45 The basis of the rings test is to assume that a good minimum (or maximum) will have points of increasing (or decreasing) magnitudes surrounding it. Figure 24 illustrates this assumption, showing a minimum at (0, 0) where the surrounding three rings of points have decreasing mean magnitude. This is as opposed to the correlation surface shown in Figure 25, where the rings, and in particular the second inner-most ring, are not of decreasing mean magnitude.

50 In this case the criteria for a good minimum as defined by the rings test, is that the average slope is monotonic. Therefore for a pre-defined number of rings of points surrounding the minimum in question, the mean magnitude of each ring when moving from the innermost ring outwards, must be greater than that of the previous ring. Returning again to the example described with reference to Figures 18 to 21, it will be seen from Figures 20 and 21 that the correlation surface of Figure 20 would have failed the rings test, but 55 that the correlation surface of Figure 21 would have passed the rings test. Since the rings test compares mean, and not absolute, magnitudes, it is far less subjective than the threshold test, and indeed the only variable in the rings test is the number of rings considered.

Having described the mechanism for growing a search block, it is now necessary to consider how by

examining the shape of the correlation surface it is possible to determine the most effective direction in which the search block should grow.

Referring again to Figure 20, this correlation surface resulted where there was a unique vertical feature, but no unique horizontal feature. This is mirrored in the correlation surface by the minimum running horizontally across the correlation surface, due to the multiple correlations in this direction. From this it can be deduced that the search block should be grown horizontally. Conversely, should a line of multiple correlations run vertically, this would indicate the need to grow the search block vertically, whilst a circular collection of multiple correlations would indicate a need to grow the search block both horizontally and vertically.

Using this criteria, a quantitative measure of the shape of the correlation surface is required in order to determine in which direction the search block should be grown. This measure is determined as follows. Firstly, a threshold is determined. Any point on the correlation surface below the threshold is then considered. This threshold, like that used in the threshold test, is normalized as a fraction of the range of magnitudes within the correlation surface. Using this threshold, the points on the correlation surface are examined in turn in four specific sequences. In each, the point at which the correlation surface value falls below the threshold is noted. These four sequences are illustrated diagrammatically in Figure 26 in which the numbers 1, 2, 3 and 4 at the top, bottom, left and right refer to the four sequences, and the hatched area indicates points which fall below the threshold:

Sequence 1

Search from the top of the correlation surface down for a point A which falls below the threshold.

Sequence 2

Search from the bottom of the correlation surface up for a point C which falls below the threshold.

Sequence 3

Search from the left of the correlation surface to the right for a point D which falls below the threshold.

Sequence 4

Search from the right of the correlation surface to the left for a point B which falls below the threshold.

The locations of the four resulting points A, B, C and D are used to calculate the two dimensions X and Y indicated in Figure 26, these dimensions X and Y indicating the size of the hatched area containing the points falling below the threshold value. Hence from the dimensions X and Y, it can be deduced whether the shape is longer in the x rather than the y direction, or vice versa, or whether the shape is approximately circular. A marginal difference of say ten percent is allowed in deducing the shape, that is, the dimension X must be a minimum of ten percent greater than the dimension Y for the shape to be considered to be longer in the x direction. Similarly for the y direction. If the dimensions X and Y are within ten percent of each other, then the shape is considered to be circular, and the search block is grown in both directions. In the example of Figure 26 the dimension X is greater than the dimension Y, and hence the search block is grown in the x or horizontal direction.

The growing of the search block continues until one or more growth limitations is reached. These limitations are: that the minimum in the correlation surface passes both the threshold test and the rings test; that the edge of the video frame is reached; or that the search block has already been grown a predetermined number of times horizontally and vertically. This last limitation is hardware dependent. That is to say, it is limited by the amount of processing that can be done in the available time. In one specific embodiment of apparatus according to the present invention, this limit was set at twice horizontally and once vertically.

If the minimum in the correlation surface passes both the threshold test and the rings test, then it is assumed that a good motion vector has been determined, and can be passed to the motion vector reducer 5 (Figure 1). However, if the edge of the frame is reached or the search block has already been grown a predetermined number of times both horizontally and vertically, then it is assumed that a good motion vector has not been determined for that particular search block, and instead of attempting to determine a good motion vector, the best available motion vector is determined by weighting.

The correlation surface is weighted such that the selection of the best available motion vector is weighted towards the stationary, that is the centre, motion vector. This is for two reasons, firstly, if the search block, even after growing, is part of a large plain area of source video, it will not be possible to detect a good motion vector. However, since the source video is of a plain area, a stationary motion vector will lead to the correct results in the subsequent processing. Secondly, weighting is designed to reduce the possibility of a seriously wrong motion vector being passed to the motion vector reducer 5 (Figure 1). This is done because it is assumed that when a good motion vector cannot be determined, a small incorrect motion vector is preferable to a large incorrect motion vector.

Figure 27 shows an example of how the weighting function can be applied to the correlation surface. In

this example, the weight applied to a given point on the correlation surface is directly proportional to the distance of that point from the stationary, centre motion vector. The magnitude of the point on the correlation surface is multiplied by the weighting factor. For example, the gradient of the weighting function may be such that points plus or minus 32 pixels from the centre, stationary motion vector are multiplied by a factor of three. In other words, as shown in Figure 27, where the centre, stationary motion vector is indicated by the black circle, the weighting function is an inverted cone which is centred on the centre, stationary motion vector.

After the correlation surface has been weighted, it is again passed through the threshold test and the rings test. If a minimum which passes both these tests is determined, then it is assumed that this is a good motion vector, and it is flagged to indicate that it is a good motion vector, but that weighting was used. This flag is passed, together with the motion vector to the motion vector reducer 5 (Figure 1). If on the other hand, neither a good motion vector nor a best available motion vector can be determined, even after weighting, then a flag is set to indicate that any motion vector passed to the motion vector reducer 5 (Figure 1) for this search block is a bad motion vector. It is necessary to do this because bad motion vectors must not be used in the motion vector reduction process, but must be substituted as will be described below.

Thus, in summary, the operation of the motion vector estimator 4 (Figure 1) is to derive from the correlation surface generated by the direct block matcher 3 (Figure 1), the point of best correlation, that is the minimum. This minimum is then subjected to the threshold test and the rings test, both of which the minimum must pass in order for it to be considered to represent the motion of the search block. It should, incidentally, be noted that the threshold used in the threshold test and the rings test may be either absolute values or fractional values. If the minimum fails either test, then the search block is grown, a new minimum is determined, and the threshold test and the rings test re-applied. The most effective direction in which to grow the search block is determined from the shape of the correlation surface.

Referring initially to Figure 1, the process of motion vector reduction will now be described. Using a HDVS, each search block is assumed to be 32 pixels by 23 lines, which can be shown to lead to a possible maximum of 2451 motion vectors. The choice of the search block size is a compromise between maintaining resolution and avoiding an excessive amount of hardware. If all these motion vectors were passed to the motion vector selector 6, the task of motion vector selection would not be practicable, due to the amount of processing that would be required. To overcome this problem, the motion vector reducer 5 is provided between the motion vector estimator 4 and the motion vector selector 6. The motion vector reducer 5 takes the motion vectors that have been generated by the motion vector estimator 4 and presents the motion vector selector 6 with only, for example, four motion vectors for each search block in the frame, including those in border regions, rather than all the motion vectors derived for that frame. The effect of this is two-fold. Firstly, this makes it much easier to choose the correct motion vector, so long as it is within the group of four motion vectors passed to the motion vector selector 6. Secondly, however, it also means that if the correct motion vector is not passed as one of the four, then the motion vector selector 6 is not able to select the correct one. It is therefore necessary to try to ensure that the motion vector reducer 5 includes the correct motion vector amongst those passed to the motion vector selector 6. It should also be mentioned that although four motion vectors are passed by the motion vector reducer 5 to the motion vector selector 6, only three of these actually represent motion, the fourth motion vector always being the stationary motion vector which is included to ensure that the motion vector selector 6 is not forced into applying a motion vector representing motion to a stationary pixel. Other numbers of motion vectors can be passed to the motion vector selector 6, for example, in an alternative embodiment four motion vectors representing motion and the stationary motion vector may be passed.

Hereinafter the term 'sample block' refers to a block in a frame of video in which each pixel is offered the same four motion vectors by the motion vector reducer 5. Thus, a sample block is the same as a search block before the search block has been grown. As shown in Figure 28, in a frame of video the initial positions of the sample blocks and the search blocks are the same.

The motion vector reducer 5 (Figure 1) receives the motion vectors and the flags from the motion vector estimator 4 (Figure 1) and determines the quality of the motion vectors by examining the flags. If the motion vector was not derived from an ambiguous surface, that is there is a high degree of confidence in it, then it is termed a good motion vector, but if a certain amount of ambiguity exists, then the motion vector is termed a bad motion vector. In the motion vector reduction process, all motion vectors classed as bad motion vectors are ignored, because it is important that no incorrect motion vectors are ever passed to the motion vector selector 6 (Figure 1), in case a bad motion vector is selected thereby. Such selection would generally result in a spurious dot in the final picture, which would be highly visible.

Each of the motion vectors supplied to the motion vector reducer (Figure 1) was obtained from a

particular search block, and hence a particular sample block (Figure 28), the position of these being noted together with the motion vector. Because any motion vectors which have been classed as bad motion vectors are ignored, not all sample blocks will have a motion vector derived from the search block at that position. The motion vectors which have been classed as good motion vectors, and which relate to a particular search block, and hence a particular sample block, are called local motion vectors, because they have been derived in the area from which the sample block was obtained. In addition to this, another motion vector reduction process counts the frequency at which each good motion vector occurs, with no account taken of the actual positions of the search blocks that were used to derive them. These motion vectors are then ranked in order of decreasing frequency, and are called common motion vectors. In the worst case only three common motion vectors are available and these are combined with the stationary motion vector to make up the four motion vectors to be passed to the motion vector selector 6 (Figure 1). However, as there are often more than three common motion vectors, the number has to be reduced to form a reduced set of common motion vectors referred to as global motion vectors.

A simple way of reducing the number of common motion vectors is to use the three most frequent common motion vectors and disregard the remainder. However, the three most frequent common motion vectors are often those three motion vectors which were initially within plus or minus one pixel motion of each other vertically and/or horizontally. In other words, these common motion vectors were all tracking the same motion with slight differences between them, and the other common motion vectors, which would have been disregarded, were actually tracking different motions.

In order to select the common motion vectors which represent all or most of the motion in a scene, it is necessary to avoid choosing global motion vectors which represent the same motion. Thus, the strategy actually adopted is first to take the three most frequently occurring common motion vectors and check to see if the least frequent among them is within plus or minus one pixel motion vertically and/or plus or minus one pixel motion horizontally of either of the other two common motion vectors. If it is, then it is rejected, and the next most frequently occurring common motion vector is chosen to replace it. This process is continued for all of the most frequently occurring common motion vectors until there are either three common motion vectors which are not similar to each other, or until there are three or less common motion vectors left. However, if there are more than three common motion vectors left, then the process is repeated this time checking to see if the least frequent among them is within plus or minus two pixel motion vertically and/or plus or minus two pixel motion horizontally of another, and so on at increasing distances if necessary. These three common motion vectors are the required global motion vectors, and it is important to note that they are still ranked in order of frequency.

When considering the motion vector reduction process and the sample blocks of a frame of video, it is necessary to look at three different types of sample blocks. These types are related to their actual position in a frame of video, and are shown in Figure 29 as regions. Region A comprises sample blocks which are totally surrounded by other sample blocks and are not near the picture boundary. Region B contains sample blocks which are partially surrounded by other sample blocks and are not near the picture boundary. Finally, region C contains sample blocks which are near the picture boundary. The motion vector reduction algorithm to be used for each of these regions is different. These algorithms will be described below, but firstly it should be reiterated that there exist good motion vectors for some of the sample blocks in the frame of video, and additionally there are also three global motion vectors which should represent most of the predominant motion in the scene. A selection of these motion vectors is used to pass on three motion vectors together with the stationary motion vector for each sample block.

Figure 30 illustrates diagrammatically motion vector reduction in the region A. This is the most complex region to deal with, because it has the largest number of motion vectors to check. Figure 30 shows a central sample block which is hatched, surrounded by other sample blocks a to h. Firstly, the locally derived motion vector is examined to see if it was classed as a good motion vector. If it was, and it is also not the same as the stationary motion vector, then it is passed on. However, if it fails either of these tests, it is ignored. Then the motion vector associated with the sample block d is checked to see if it was classed as a good motion vector. If it was, and if it is neither the same as any motion vector already selected, nor the same as the stationary motion vector, then it too is passed on. If it fails any of these tests then it too is ignored. This process then continues in a similar manner in the order e, b, g, a, h, c and f. As soon as three motion vectors, not including the stationary motion vector, have been obtained, then the algorithm stops, because that is all that is required for motion vector selection for that sample block. It is, however, possible for all the above checks to be carried out without three good motion vectors having been obtained. If this is the case, then the remaining spaces are filled with the global motion vectors, with priority being given to the more frequent global motion vectors.

Figure 31 illustrates motion vector reduction in the region B. Sample blocks in the region B are the

same as those in the region A, except that they are not totally surrounded by other sample blocks. Thus the process applied to these sample blocks is exactly the same as those for the region A, except that it is not possible to search in all the surrounding sample blocks. Thus as seen in Figure 31, it is only possible to check the motion vectors for the sample blocks a to e, and any remaining spaces for motion vectors are filled, as before, with global motion vectors. Likewise, if the hatched sample block in Figure 31 were displaced two positions to the left, then it will be seen that there would only be three adjacent surrounding blocks to be checked before resorting to global motion vectors.

Figure 32 illustrates motion vector reduction in the region C. This is the most severe case, because the sample blocks neither have a locally derived motion vector nor do they have many surrounding sample blocks whose motion vectors could be used. The simplest way of dealing with this problem is simply to give the sample blocks in the region C the global motion vectors together with the stationary motion vector. However, this is found to produce a block-like effect in the resulting picture, due to the sudden change in the motion vectors presented for the sample blocks in the region C compared with adjoining sample blocks in the region B. Therefore a preferred strategy is to use for the sample blocks in the region C the sample motion vectors as those used for sample blocks in the region B, as this prevents sudden changes. Preferably, each sample block in the region C is assigned the same motion vectors as that sample block in the region B which is physically nearest to it. Thus, in the example of Figure 32, each of the hatched sample blocks in the region C would be assigned the same motion vectors as the sample block a in the region B, and this has been found to give excellent results.

This assignment of motion vectors to the sample blocks in the region B, with which the present invention is particularly, but not exclusively, concerned, will now be described in more detail with reference to Figures 33 to 35.

Figure 33 shows a circuit arrangement for effecting the required assignment in the region B. The arrangement comprises a count system including counters 21 to 24, buffers 25, 26, 28, 29, 32 and 33, a microprocessor 27, a random access memory (RAM) 30, and a read only memory (ROM) 31 which holds the program for the microprocessor 27, these elements being interconnected as shown and operating as will now be described.

The count system 20 is reset at the start of a frame and then the motion vectors for that frame are loaded by way of the buffer 32. While the motion vectors are being loaded, the microprocessor 27 is disabled, as it is also while the motion vectors are subsequently being read out by way of the buffer 33.

While the motion vectors are being loaded, the counter 21 counts the input motion vectors, and hence controls the input of the motion vectors to the RAM 30 by providing sample block addresses by way of the buffer 25. Meanwhile, the counter 23 counts the sample blocks in the horizontal direction of the frame, and the counter 24 counts the sample blocks in the vertical direction of the frame. On completion of the frame, the counter 24 causes the counter 22 to control the output of the motion vectors from the RAM 30 by providing sample block addresses by way of the buffer 26.

The RAM 30 has three areas. A first area is set aside to store the original motion vectors relating to the sample blocks. This area will only be read by the microprocessor 27, and will never be written in by the microprocessor 27. A second area is set aside to store the output motion vectors, as this makes reading out easier. A third area is set aside for processing.

Initially then the microprocessor 27 and the motion vector output section are disabled while the incoming motion vectors are written in the RAM 30. The motion vector input and output sections are then disabled while processing according to the algorithm described below is performed. Then the motion vector input section and the microprocessor 27 are disabled while the motion vectors are read out of the RAM 30 by the motion vector output section by way of the buffer 33.

Referring to Figure 34, this shows a frame of video, and in particular sample blocks corresponding to the region B, that is the sample blocks in the region immediately adjacent to but not actually in the border region C.

First of all consider the sample block a. If the local motion vector that has been generated for the sample block a is both good and non-zero, then it is retained as one of the four motion vectors for the sample block a. Next, the local motion vector generated for the sample block d is compared with that already retained for the sample block a. If it is the same as the local motion vector for the sample block a or is zero then it is discarded. However, if it is good, non-zero and not the same as the local motion vector for the sample block a then it too is retained. Then the local motion vector for the sample block b is compared with those already retained for the sample block a. If it is the same as a motion vector already retained or is zero it is discarded. However, if it is good, non-zero and not the same as a motion vector already retained it too is retained. Then the local motion vector for the sample block c is examined in a similar way, but in this case against the local motion vectors for the sample blocks a, d and b if they have been retained.

The remainder of the sample blocks adjacent to the border region C are motion vector reduced in a similar way. However, the sample blocks used for the selection of motion vectors to be retained depend on the location of the original sample block relative to the border region C. The following lists examples of the various different locations that the algorithm has to deal with, and in each case shows the sample blocks considered, and the order of this consideration.

- Sample block a : sample blocks d, b, c
- Sample block e : sample blocks f, h, g
- Sample block q : sample blocks r, f, s, t, g
- Sample block u : sample blocks v, h, w, g, x
- 10 Sample block l : sample blocks l, j, k
- Sample block i : sample blocks y, i, k, z, j
- Sample block m : sample blocks p, n, o

Motion vector reduction stops if either three motion vectors have been retained, or there are no sample blocks left to be examined. If three motion vectors have not been retained for a sample block, then the 15 remaining spaces will be filled by using the global motion vectors in order of frequency. In this way each sample block should get three different non-zero motion vectors plus the zero, that is the stationary motion vector which is always supplied.

The operation of the arrangement of Figure 33 in accordance with this algorithm is shown in the form of a flow chart in Figure 35.

20 Referring again to Figure 1, the purpose of the motion vector selector 6 is to assign one of the four motion vectors supplied thereto to each individual pixel within the sample block. In this way the motion vectors can be correctly mapped to the outline of objects. The way in which this assignment is effected is particularly intended to avoid the possibility of the background surrounding fine detail from producing a better match than that produced by the correct motion vector. To achieve this the motion vector selection 25 process is split into two main stages. In the first stage, motion vectors are produced for each pixel in the input frames. In other words, there is no attempt to determine the motion vector values for pixels at the output frame positions. The second stage uses the motion vector values produced by the first stage to determine the motion vector value for each pixel in the output frame.

Referring now to Figure 33, each pixel of the input frame 2 is tested for the best luminance value match 30 with the previous and following input frames 1 and 3 of video data, using each of the four motion vectors supplied. The pixel luminance difference is determined as:

$$35 \quad \sum_{m=0}^4 \sum_{n=0}^4 |P1_{nm} - P2_{nm}| + \sum_{m=0}^4 \sum_{n=0}^4 |P2_{nm} - P3_{nm}|$$

where:

40  $P1_{nm}$  is the luminance value of a frame 1 pixel within a 4x4 block of pixels surrounding the pixel whose location is obtained by subtracting the coordinate of the motion vector being tested from the location of the pixel being tested in frame 2

$P2_{nm}$  is the luminance value of a frame 2 pixel within a 4x4 block of pixels surrounding the pixel being tested

45  $P3_{nm}$  is the luminance value of a frame 3 pixel within a 4x4 block of pixels surrounding the pixel whose location is obtained by adding the coordinates of the motion vector being tested to the location of the pixel being tested in frame 2

The minimum pixel difference then indicates the best luminance match and therefore the correct motion 50 vector applicable to the pixel being tested. If the correct motion vector is not available, or there are uncovered or covered areas, referred to in more detail below, then a good match may not occur.

The indication of a poor match is achieved when the average pixel difference within the block of pixels being used is above a certain threshold. This threshold is important, because high frequency detail may produce a poor match even when the correct motion vector is tested. The reason for this poor match is the possibility of a half pixel error in the motion vector estimate. To determine what threshold should indicate a poor match, it is necessary to relate the threshold to the frequency content of the picture within the block of data which surrounds the pixel for which the motion vector is required. To achieve this, an auto-threshold value is determined where the threshold value equals half the maximum horizontal or vertical pixel luminance difference about the pixel being tested. To ensure that the threshold value obtained is

representative of the whole block of data which is compared, an average value is obtained for the four central pixels of a 4x4 block used.

Referring to Figure 38, which shows a 4x4 block, the required threshold value  $T$  is given by:

5  $T = (T_1 + T_2 + T_3 + T_4)/8$   
 where  $T_3$ , for example, is determined as indicated in Figure 39 as equal to the maximum of the four pixel  
 10 luminance difference values comprising:  
 the two vertical differences  $|B_2 - B_3|$  and  $|B_4 - B_3|$ , and  
 the two horizontal differences  $|A_3 - B_3|$  and  $|C_3 - B_3|$ .

15 In this way a frame of motion vectors is obtained for input frame 2, and in a similar manner a frame of motion vectors is obtained for input frame 3 as indicated in Figure 37.

16 Apart from scene changes, it is the phenomenon of uncovered/covered surfaces that causes a mismatch to occur in the above first stage of motion vector selection. If an object, say a car, drives into a tunnel, then the car has become covered, while when it drives out, the car is uncovered. If the part of the car that was uncovered in frames 1 and 2 is covered in frames 3 and 4, then the basic vector selection process is not able to determine the correct vector. Moreover, whilst the car going into the tunnel becomes covered, the road and objects behind the car are being uncovered. Likewise the car leaving the tunnel is being uncovered, but the road and objects behind the car are being covered. In general therefore both covered and uncovered objects will exist at the same time. The end of a scene will also have a discontinuation of motion that is similar to an object becoming covered. In an attempt to determine a motion vector even in such circumstances, the luminance value block match is reduced to a two frame match, instead of the three frame match of Figures 36 and 37. The frame that the motion vectors are required for (say frame 2) is block-matched individually to the previous and the next frame (frame 1 and frame 3 respectively, in the case of frame 2), using the four motion vectors supplied. The motion vector which produces the best match is chosen as the motion vector applicable to the pixel being tested. In this case, 20 however, a flag is set to indicate that only a two frame match was used.

25 Particularly with integrating type television cameras, there will be situations where no match occurs. If an object moves over a detailed background, then an integrating camera will produce unique portions of picture where the leading and trailing edges of the object are mixed with the detail of the background. In such circumstances, even the two frame match could produce an average pixel difference above the threshold value. In these cases the motion vector value is set to zero, and an error flag is also set.

30 The second stage of motion vector selection makes use of the two frames of motion vectors, derived by the first stage. One frame of motion vectors (input frame 2) is considered to be the reference frame, and the following frame to this (input frame 3) is also used. The output frame position then exists somewhere between these two frames of motion vectors. Referring to Figure 40, for each output pixel position the four possible motion vectors associated with the sample block of input frame 2, are tested. A line drawn through the output pixel position at the angle of the motion vector being tested will point to a position on both the input frame 2 and the input frame 3. In the case of odd value motion vectors, for example, 1, 3 and 5, a point midway between two input frame pixels would be indicated. In the case where the output frame is precisely half way between the input frames 1 and 2. To allow for this inaccuracy, and also to reduce the sensitivity to individual pixels, a 3x3 block of motion vectors is acquired for each frame, centred on the closest pixel position. In effect a block-match is then performed between each of the two 3x3 blocks of motion vectors and a block containing the motion vector being tested. The motion vector difference used represents the spatial difference of the two motion vector values as given by:

$$\sqrt{((x_1 - x_2)^2 + (y_1 - y_2)^2)}$$

45 where:

$x_1$  and  $y_1$  are the Cartesian coordinates of the motion vector in one of the blocks

$x_2$  and  $y_2$  are the Cartesian coordinates of the motion vector being tested

An average vector difference per pixel is produced as a result of the block match.

50 A motion vector match is first produced as above using only motion vector values which were calculated using three input frames; that is, input frames 1, 2 and 3 for input frame 2 (Figure 36), and input frames 2, 3 and 4 for input frame 3 (Figure 37), and the result is scaled accordingly. Preferably there are at least four usable motion vectors in the block of nine. When both the motion vector block of frame 2 and frame 3 can be used, the motion vector difference values are made up of half the motion vector difference value from frame 2 plus half the motion vector difference value from frame 3. Whichever motion vector 55 produces the minimum motion vector difference value using the above technique is considered to be the motion vector applicable to the output pixel being tested. If the motion vector difference value produced by the three frame match input motion vector (Figures 36 and 37) is greater than unity, then a covered or uncovered surface has been detected, and the same process is repeated, but this time ignoring the error

flags. That is, the motion vector values which were calculated using two input frames are used. Theoretically this is only necessary for uncovered/covered surfaces, although in fact improvements can be obtained to the picture in more general areas.

5 If after both of the above tests have been performed, the minimum motion vector match is greater than two, the motion vector value is set to zero, and an error flag is set for use by the motion vector post processor 7 (Figure 1).

Following motion vector selection, there will almost certainly be in any real picture situation, some remaining spurious motion vectors associated with certain pixels. Figures 41 to 46 show what are taken to be spurious motion vectors, and in each of these figures the triangles represent pixels having associated therewith the same motion vectors, whilst the stars represent pixels having associated therewith motion vectors different those associated with the surrounding pixels, and the circle indicates the motion vector under test.

Figure 41 shows a point singularity where a single pixel has a motion vector different from those of all the surrounding pixels.

15 Figure 42 shows a horizontal motion vector impulse, where three horizontally aligned pixels have a motion vector different from those of the surrounding pixels.

Figure 43 shows a vertical motion vector impulse where three vertically aligned pixels have a motion vector different from those of the surrounding pixels.

20 Figure 44 shows a diagonal motion vector impulse where three diagonally aligned pixels have a motion vector different from those of all the surrounding pixels.

Figure 45 shows a horizontal plus vertical motion vector impulse, where five pixels disposed in an upright cross have a motion vector different from those of all the surrounding pixels.

Figure 46 shows a two-diagonal motion vector impulse where five pixels arranged in a diagonal cross have a motion vector different from those of all the surrounding pixels.

25 It is assumed that pixel motion vectors which fall into any of the above six categories do not actually belong to a real picture, and are a direct result in of an incorrect motion vector selection. If such motion vectors were used during the interpolation process, then they would be likely to cause dots on the final output picture, and it is therefore preferable that such motion vectors be identified and eliminated. This is done using an algorithm which will detect and flag all of the above motion vector groupings.

30 The algorithm uses a two-pass process, with each pass being identical. The need for two passes will become apparent. Figure 47, to which reference is made, shows an array of pixels, all those marked with a triangle having the same motion vector associated therewith. The block of nine pixels in the centre has motion vectors designated vector 1 to vector 9 associated therewith, which motion vectors may or may not be the same. Vector 5 is the motion vector under test.

35 In the first pass, vector 5 is checked to determine whether it is the same as, or within a predetermined tolerance of:

firstly  
vector 1 or vector 3 or vector 7 or vector 9  
and secondly

40 vector 2 or vector 4 or vector 6 or vector 8

This checks to see if vector 5 is the same as at least one of its horizontal or vertical neighbours, and the same as at least one of its diagonal neighbours. If this is not the case, then a flag is set to indicate that pixel 5 is bad.

45 The first pass will flag as bad those motion vectors relating to point singularities, horizontal motion vector impulses, vertical motion vector impulses, diagonal motion vector impulses and two diagonal motion vector impulses (Figures 41 to 44 and 46), but not the motion vectors corresponding to horizontal plus vertical motion vector impulses (Figure 45) for which pass 2 is required. The second pass checks for exactly the same conditions as in the first pass, but in this case motion vectors which have already been flagged as bad are not included in the calculation. Thus, referring to Figure 45, after the first pass only the centre motion vector is flagged as bad, but after the second pass all five of the motion vectors disposed in the upright cross are flagged as bad.

Having identified the bad motion vectors, it is then necessary to repair them, this also being effected by the motion vector post processor 7 (Figure 1). Although various methods such as interpolation or majority replacement can be used, it has been found that in practice simple replacement gives good results. This is effected as follows (and it should be noted that the 'equals' signs mean not only exactly equal to, but also being within a predetermined tolerance of): If vector 5 is flagged as bad then it is replaced with:

vector 4 if (vector 4 equals vector 6)

else with vector 2 if (vector 2 equals vector 8)

else with vector 1 if (vector 1 equals vector 9)  
else with vector 3 if (vector 3 equals vector 7)  
else do nothing

Referring again to Figure 1, the finally selected motion vector for each pixel is supplied by the motion vector post processor 7 to the interpolator 8, together with the progressive scan converted frames at 60 frames per second from the progressive scan converter 2. The interpolator 8 is of relatively simple form using only two progressive scan converted frames, as indicated in Figure 4b. Using the temporal position of the output frame relative to successive input frames, frame 1 and frame 2, and the motion vector for the pixel in the output frame, the interpolator 8 determines in known manner which part of the first frame should be combined with which part of the second frame and with what weighting to produce the correct output pixel value. In other words, the interpolator 8 adaptively interpolates along the direction of movement in dependence on the motion vectors to produce motion compensated progressive scan frames corresponding to 24 frames per second. Although the motion vectors have been derived using only luminance values of the pixels, the same motion vectors are used for deriving the required output pixel chrominance values. An 8x8 array of pixels are used from each frame to produce the required output. Thus the interpolator 8 is a two-dimensional, vertical/horizontal, interpolator and the coefficients used for the interpolator 8 may be derived using the Remez exchange algorithm which can be found fully explained in 'Theory and application of digital signal processing', Lawrence R Rabiner, Bernard Gold, Prentice-Hall Inc., pages 136 to 140 and 227.

Figure 48 shows diagrammatically the interpolation performed by the interpolator 8 (Figure 1) for three different cases. The first case, shown on the left, is where there are no uncovered or covered surfaces, the second case, shown in the centre, is where there is a covered surface, and the third case, shown on the right, is where there is an uncovered surface. In the case of a covered surface, the interpolation uses only frame 1. whilst in the case of an uncovered surface, the interpolation uses only frame 2.

25 Provision can be made in the interpolator 8 to default to non-motion compensated interpolation, in which case the temporally nearest progressive scan converted frame is used.

Attention is drawn to the fact that this application is one of a series of fourteen filed on the same day, and bearing our references E241-5 (corresponding to UK 8909843.2), E241-6 (corresponding to UK 8909844.0), E241-7 (corresponding to UK 8909845.7), E241-8 (corresponding to UK 8909846.5), E241-9 (corresponding to UK 8909847.3), E241-10 (corresponding to UK 8909848.1), E241-11 (corresponding to UK 8909849.9), E241-12 (corresponding to UK 8909850.7), E241-13 (corresponding to UK 8909851.5), E241-14 (corresponding to UK 8909852.3), E241-1 (corresponding to UK 8909853.1), E242-2 (corresponding to UK 8909854.9), E242-3 (corresponding to UK 8909855.6) and E242-4 (corresponding to UK 8909856.4). These applications all relate to similar subject matter, and the disclosure in each is incorporated by this reference into each other.

## Claims

40 1. A motion compensated video standards converter comprising:  
means (3, 4) for comparing blocks in a first field or frame of a video signal with a plurality of blocks in the  
following field or frame of the video signal for deriving motion vectors representing the motion of the content  
of respective said blocks between said first field or frame and said following field or frame;  
means (27) to assign to each of said blocks adjacent to a border region of each said field or frame, motion  
45 vectors selected from motion vectors which have been derived for those of said blocks which are adjacent  
to said block under consideration and are not in said border region; and  
an interpolator (8) controlled in dependence on said motion vector.

2. A converter according to claim 1 wherein said means (27) assigns to each of said blocks adjacent to a corner of said border region, motion vectors selected from the motion vectors which have been derived for the three said blocks which are adjacent to said block under consideration and are not in said border region.

3. A converter according to claim 1 or claim 2 wherein said means (27) assigns to each of said blocks adjacent to an edge of said border region, motion vectors selected from the motion vectors which have been derived for the five said blocks which are adjacent to said block under consideration and are not in said border region.

4. A converter according to claim 1, claim 2 or claim 3 wherein said means (27) additionally assigns a zero motion vector to said block under consideration.

5. A converter according to any one of the preceding claims wherein said comparison is effected by

comparing the sums of the luminance levels of the pixels in the respective said blocks to be compared.

8. A method of deriving motion vectors representing motion between successive fields or frames of a video signal, the method including the steps of:

5 comparing blocks in a first field or frame of the video signal with a plurality of blocks in the following field or frame of the video signal for deriving said motion vectors representing the motion of the content of respective said blocks between said first field or frame and said following field or frame; and  
 10 assigning to each of said blocks (a; q) adjacent to a border region of each said field or frame, the motion vectors selected from the motion vectors which have been derived for those of said blocks (d, b, c; r, f, s, t, g) which are adjacent to said block (a, q) under consideration and are not in said border region.

10 7. A method according to claim 6 wherein said assignment comprises assigning to each of said blocks (a) adjacent to a corner of said border region, motion vectors selected from the motion vectors which have been derived for the three said blocks (d, b, c) which are adjacent to said block (a) under consideration and are not in said border region.

15 8. A method according to claim 6 or claim 7 wherein said assignment comprises assigning to each of said blocks (q) adjacent to an edge of said border region, motion vectors selected from the motion vectors which have been derived for the five said blocks (r, f, s, t, g) which are adjacent to said block (q) under consideration and are not in said border region.

9. A method according to claim 6, claim 7 or claim 8 comprising additionally assigning a zero motion vector to said block (a, q) under consideration.

20 10. A method according to any one of claims 6 to 9 wherein said comparison is effected by comparing the sums of the luminance levels of the pixels in the respective said blocks to be compared.

25

30

35

40

45

50

55




FIG. 1.




FIG. 2.




FIG. 3.




FIG. 4.




FIG. 5.








FIG. 6.

FIG. 7.










FIG. 11.



FRAME N



FRAME (N+1)

FIG. 13.




FIG.14.




FIG. 15.




FIG. 16.




FIG. 17.






FIG. 18.

SEARCH BLOCKS.  
1A 2A 3A



FRAME [t+1]






FIG. 19.




FIG. 20.




FIG. 21.




FIG. 22.

FIG. 23.




FIG. 24.




FIG. 25.




FIG. 26.




FIG. 27.




FIG. 28 FRAME OF VIDEO




FIG. 29.




FIG. 30.




FIG. 31.





FIG. 32.






FIG. 33.

FIG. 34.




FIG. 35.




FIG.36.





FIG.37.







FIG. 40.




FIG. 41.




FIG. 42.




FIG. 43.




FIG. 44.




FIG. 45.




FIG. 46.




FIG. 47.





FIG. 48.

NO UNCOVERED OR  
COVERED SURFACES

COVERED SURFACE

UNCOVERED SURFACE







EUROPEAN SEARCH  
REPORT

EP 90 30 3975

DOCUMENTS CONSIDERED TO BE RELEVANT

Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int. Cls)		
A	EP-A-0 294 955 (SONY CORPORATION) * column 8, lines 46 - 55 * - - -	1-10	G 06 F 15/70 H 04 N 5/14 H 04 N 7/01 H 04 N 7/137 H 04 N 5/44		
A	EP-A-0 309 251 (BRITISH TELECOMMUNICATIONS PUBLIC LIMITED COMPANY) * page 4, lines 28 - 38; figures 4, 6 * - - -	1,6			
<hr/>					
<hr/>					
<p>TECHNICAL FIELDS SEARCHED (Int. Cls)</p>					
<p>G 06 F H 04 N</p>					
<hr/>					
<p>The present search report has been drawn up for all claims</p>					
Place of search	Date of completion of search	Examiner			
Berlin	07 November 91	MATERNE A G			
<hr/>					
<p>CATEGORY OF CITED DOCUMENTS</p>					
<p>X: particularly relevant if taken alone Y: particularly relevant if combined with another document of the same category A: technological background D: derivative document P: intermediate document T: theory or principle underlying the invention</p>					
<p>E: earlier patent document, but published on, or after the filing date D: document cited in the application L: document cited for other reasons</p>					
<p>G: member of the same patent family, corresponding document</p>					